

A new rule for almost-certain termination of probabilistic- and demonic programs

Annabelle McIver¹ and Carroll Morgan²

¹ Macquarie University, Australia

² University of New South Wales, and Data61, Australia.

Abstract. Extending our own and others’ earlier approaches to reasoning about termination of probabilistic programs, we propose and prove a new rule for termination with probability one, also known as “almost-certain termination”. The rule uses both (non-strict) super martingales and guarantees of progress, together, and it seems to cover significant cases that earlier methods do not. In particular, it suffices for termination of the unbounded symmetric random walk in both one- and two dimensions: for the first, we give a proof; for the second, we use a theorem of Foster to argue that a proof exists.

Non-determinism (i.e. demonic choice) is supported; but we do currently restrict to discrete distributions.³

1 Introduction

This paper concerns proof of almost-sure termination for probabilistic- and demonic programs, ones that move from one state to another by first choosing a (discrete) distribution demonically from a set of distributions, and then choosing a new state probabilistically according to that distribution.

Thus we view a program abstractly as a (probabilistic/demonic) transition system; and we are interested in proving the eventual reachability with probability one of a given set of target states (when it is indeed the case). Our strategic aim is to express our techniques in a form that can be applied to probabilistic program code *in situ*, i.e. without the need to construct the programs’ underlying transition systems explicitly (although of course we rely on their existence). That is, we seek (and find) proof-rules that require no more than local reasoning in the source code.

In probabilistic programming over a finite state-space S , say, a typical rule is one that generalises the “variant rule” for standard (non-probabilistic but still demonic) programs: to each state is assigned a non-negative integer, a *variant* bounded above and below, with all states inside some target set $S_0 \subseteq S$ assigned variant 0. One then shows by local reasoning, typically over the source-code of

³ This revises an earlier version [22] by correcting typographical errors and adding an extra section §9 on historical background.

a loop body,

that from each state in $S_* = S - S_0$ there is a *non-zero probability* of transiting to a (different) state with strictly smaller (integer) variant. ⁴ (1)

If that can be shown, then the rule guarantees that almost all paths in the transition system eventually lead to S_0 , where “almost all” means that the paths not included have probability zero even if taken all together. This probabilistic rule’s soundness follows from an appeal to a zero-one law [12, 23, 20] which roughly says:

If there is some $\varepsilon > 0$ such that the probability of eventually reaching a target set of states is *everywhere* at least ε , then that probability is one.

For infinite-state systems however, although such zero-one laws are still valid, their ε -conditions are not so easily met by local reasoning. In particular the actual values of the probabilities attached to the transitions, which in fact are irrelevant in finite-state transition systems [24, 21], now can make the difference between almost-certain termination or not. A typical response to this issue is to replace “non-zero probability” in (1) above with “probability bounded away from zero”. And that bound can depend intricately on the transitions’ actual probabilities.

That challenge notwithstanding, recent important work [4, 8, 3] has shown how local reasoning with super-martingales can be applied to solve the termination problem in a wide class of infinite-state probabilistic programs.

In this paper we combine those successes with some of our own earlier work, showing in this paper how to use super-martingale reasoning together with a progress rule to reason about to an important class of transition systems whose termination seems to be beyond the state-of-the-art, for source-level reasoning at least. Our key insight is the observation that the combination of a super-martingale with a local *but parametrised* progress-condition (in a sense we explain below) implies the conditions of the zero-one rule.

Our specific contributions are:

1. A new rule that generalises a number of currently known rules (including our own) for establishing almost-certain termination;
2. A demonstration of a general zero-one proof technique which can be applied in arbitrary infinite state systems;
3. A thorough analysis of the applicability of the new rule together with a suite of representative examples; and
4. A limited survey of some pre- computer-science mathematical results that contribute to this endeavour [16, 9, 1]

⁴ This is the probabilistic generalisation: in the traditional, non-probabilistic rule the decrease must be certain. On the other hand, in the traditional case the variant need not be bounded above. In both cases, it must be bounded below.

Finally we note that our strategic goal, to translate this and other rules to ones that can be applied directly to program code, lies in the seminal work of Kozen [18] for probabilistic semantics, later generalised by us to include demonic non-determinism and abstract transition systems [25, 20] and even more recently expanded to include explicit Markov-chain models [10].

2 Informal description of the new rule for termination

2.1 Setting of the new rule, and its purpose

Let S be a state space, possibly infinite, and let S be divided into two disjoint subsets: one is S_0 , the states where termination is deemed to have occurred; and the other is S_* for the rest. A transition function T is given taking any state s in S_* to a set of discrete distributions on all of S ; and a transition from some s in S_* occurs by first selecting arbitrarily a distribution δ from $T(s)$ and then choosing probabilistically a next-state s' according to the probabilities given by δ . (We discuss this treatment of demonic nondeterminism more thoroughly in §3 below.)

Our purpose is to give a method for proving that from any s in S_* , repeated transitions according to T will reach S_0 with probability one eventually. That property is conventionally called *almost-certain termination*. For brevity, from now on we will write *AC* for “almost certain(ly)” and *ACT* for “almost-certain(ly) terminat(ion/ing)”.

2.2 Informal description of the rule

By analogy with existing approaches to proof of termination, we base our technique on a “variant” function over the states and require it to have certain properties. Informally described, they are as follows:

Define a *variant* — Non-negative variant function V from (all of) S into the non-negative reals is such that V is 0 on all of S_0 and V is strictly positive on S_* . It can be unbounded (above), but not infinite. Note that V need not be integer-valued.

Impose a *super-martingale* property — Variant V is a super-martingale wrt. transitions T , i.e. for any s in S_* and any distribution δ in $T(s)$, the expected value of V on δ , i.e. on the states reached in one δ -mediated transition of T from s , is no more than the value $V(s)$ that V had at s itself. That is,

$$\text{For all } s \text{ in } S_* \text{ and } \delta \text{ in } T(s) \text{ we have } \mathcal{E}_\delta V \leq V(s) \quad ,$$

where we write $\mathcal{E}_\delta V$ for the expected value, over discrete distribution δ on S , of real-valued function V on S .

Note that we do not require a strict decrease of the expected value, and although V is defined on S_0 , we do not require that T be defined there.⁵

⁵ Although in general there is a question of definedness of $\mathcal{E}_\delta V$ when δ has infinite support and V is unbounded, that does not arise here.

Impose a *progress* property — The transitions T makes progress towards S_0 . We require two fixed strictly positive functions p (for “probability”) and d (for “decrease”), defined for all positive reals, such that in a state s of S_* with $V(s)$ equal to some v , any transition δ in $T(s)$ is guaranteed with probability at least $p(v)$ to decrease the variant by at least $d(v)$. Furthermore $d(v)$ and $p(v)$ must be non-increasing as v itself increases. That is,

There are fixed functions p, d on positive reals v , with $0 < p(v) \leq 1$ and $0 < d(v)$, such that whenever $v = V(s)$ for some s in S_* , and δ in $T(s)$, we have

$$\delta_{\{s' \mid V(s') \leq v - d(v)\}} \geq p(v) , \quad \text{where for } S' \subseteq S \\ \text{we write } \delta_{S'} \text{ for } \sum_{s' \in S'} \delta_{s'} ,$$

and for any $0 < v < v'$ we have $p(v') \leq p(v)$ and $d(v') \leq d(v)$.

Note that p, d in *progress* are functions of the *variant*, defined over all positive reals, and that even for v not in the V -image of S , still the non-increasing conditions for $p(v), d(v)$ must be satisfied. (See §4.2’s “What happens when V is bounded”.)

The rule is proved in §4.1 below.

2.3 Discussion and comparison

Our main innovation in §2.2 is, in our *progress* condition, to impose the usual “bounded away from zero” criterion not on S_* as a whole but instead only on successively larger subsets of it. That is, we apply it with respect to certain functions p and d , and the effect of their non-increasing criteria is to ensure that, as the subsets $\{s : S_* \mid V(s) \leq v\}$ of S_* grow larger, the progress conditions imposed on them grow weaker but never decrease to “none”. This avoids the treacherous Zeno-effects that can occur when some progress is always made but only with ever-smaller steps: the V -decrease condition (“as far as d with probability at least p ”) can only be strengthened as $V(s)$ moves towards 0. But it also avoids the need to set a uniform ε -progress condition for all of S_* .

Although the generality of p, d might seem complicated, in fact in many special cases it is very simple. One such is the “distance from 0” variant on the one-dimensional symmetric random walk, where p, d can be constant functions: we take S to be the integers, both positive and negative, with $S_0 = \{0\}$ and $V(s) = |s|$ and we define p, d to be everywhere $1/2, 1$ respectively — with probability at least $1/2$ the variant decreases by at least 1. That is all that’s needed to establish *ACT* for the symmetric random walk (§7.1).⁶

⁶ This simplicity shows that the difficulty of finding an *ACT* rule lies in part in making sure it does not allow too much: what prevents our rule’s proving that a biased random walk is *ACT*? See §7.2.

2.4 Other approaches

In **our own, earlier probabilistic-variant rule** [23, Sec. 6],[20, Sec. 2.7], we effectively made p, d constants, imposed no *super-martingale* condition but instead bounded V above over S_* , making it not sufficient for the random walk. **Later however** we did prove random walk to be *ACT* using a rule more like the current one [20, Sec. 3.3].

Chakarov and Sankaranarayanan [4] consider the use of martingales for the analysis of infinite-state probabilistic programs, and **Chakarov** has done more extensive work [3].

In [4] it's shown that a *ranking super-martingale* implies *ACT*, and a key property of their definition for ranking super-martingale is that there is some constant $\varepsilon > 0$ such that the average decrease of the super-martingale is everywhere (except for the termination states) at least ε . Their program model is assumed to operate over discrete state spaces, without nondeterminism.

That work is an important step towards applying results from probability theory to the verification of infinite-state probabilistic programs.

Fioriti and Hermanns [8] also use ranking super-martingales, with results that provide a significant extension to Chakarov and Sankaranarayanan's work [4]. Their program model includes both non-determinism and continuous probability distributions over transitions. They also show completeness for the class of programs whose expected time to termination is finite. That excludes the random walk however; but they do demonstrate by example that the method can still apply to some systems which do not have finite termination time.

More recently still, **Chatterjee, Novotný and Žikelić** [5] study techniques for proving that programs terminate with some probability (not necessarily one). Their innovation is to introduce the concept of “repulsing super-martingales” — these are also super-martingales with values that decrease outside of some defined set. Repulsing super-martingales can be used to show lower bounds on termination probabilities, and as certificates to refute almost-sure termination and finite expected times to termination. (See also §9.2, §9.5.)

There are a number of other works that demonstrate tool support based on the above and similar techniques. All the authors above [4, 8, 5] have developed and implemented algorithms to support verification based on super-martingales. **Esparza, Gaiser and Kiefer** [7] develop algorithmic support for *ACT* of “weakly finite” programs, where a program is *weakly finite* if the set of states reachable from any initial state is finite. **Kaminski et al.** [15] have studied the analysis of expected termination times of infinite state systems using probabilistic invariant-style reasoning, with some applications to *ACT*. In even earlier work **Celiku and McIver** [2] explore the mechanisation of upper bounds on expected termination times, taking probabilistic weakest pre-conditions [20] for their model of probability and non-determinism.

3 Our treatment of demonic nondeterminism

Before proving §2.2, we explain our treatment of demonic- and probabilistic choice together.

Our transition function T is of type $S_* \rightarrow \mathbb{P}\mathbb{D}S$, where \mathbb{P} is the powerset constructor and \mathbb{D} is the discrete-distribution constructor: thus for state s in S_* , its possible transitions $T(s)$ comprise a set (\mathbb{P}) of discrete distributions (\mathbb{D}) of states (S). It simultaneously extends (1) the conventional model $S \rightarrow \mathbb{P}S$ of demonic (non-probabilistic) programs and e.g. (2) Kozen’s model $S \rightarrow \mathbb{D}S$ [18] and later Plotkin and Jones’ model [14] of probabilistic (non-demonic, i.e. deterministic) programs. For (1) the embedding $\mathbb{P}S \hookrightarrow \mathbb{P}\mathbb{D}S$ is as sets of point distributions, and for (2) the embedding $\mathbb{D}S \hookrightarrow \mathbb{P}\mathbb{D}S$ is as singleton sets of distributions.

The full probabilistic/demonic model has been thoroughly explored in earlier work [25, 13, 20] and has an associated simple programming language $pGCL$, for which it provides a denotational semantics.⁷

Using $pGCL$ semantics, we can model our system as a while-loop of the form

while $s \notin S_0$ **do** “choose s' according to $T(s)$ ”; $s := s'$ **end** ,

where “choose s' according to $T(s)$ ” is simply a $pGCL$ probabilistic/demonic assignment statement and the semantics of **while** is given as usual by a least fixed-point.

An alternative, more recent approach is concerned with expected time to termination, and while-loops’ semantics are given equivalently as limits of sequences of distributions [15]. Either way, the resulting set of final distributions (non-singleton, if there is nondeterminism) comprises *sub*-distributions, summing to no more than one (rather than to one exactly), where the “one deficit” is the probability of never escaping the loop. Proving *ACT* then amounts to showing that all those sub-distributions are in fact full distributions, summing to one.

Our relying on well established semantics for demonic choice and probability together is the reason we do not have to construct a scheduler explicitly, as some approaches do: the scheduler’s actions are “built in” to the set-of-distributions semantics.

4 Proof of the new rule for almost-certain termination

4.1 Proof of soundness

Theorem 1. *Soundness of §2.2* The technique described in §2.2 is sound.

⁷ This approach is also similar to the work of Segala [26], whose construction based on I/O automata appeared at about the same time as the workshop version of [13]; and it has numerous connections with probabilistic/demonic process algebras as labelled transition systems that alternate between demonic- and probabilistic branching.

Proof Recall that the state space is S , that the termination subset is $S_0 \subseteq S$ and that $S_* = S - S_0$ is the rest. The transition function T is of type $S_* \rightarrow \mathbb{P}\mathbb{D}S$ and the variant V is of type $S \rightarrow \mathbb{R}^{\geq}$ with $V(S_0) = \{0\}$.

Fix some non-negative real number H (for “high”), and consider the subset S_H of S_* whose variants are no more than H , that is $\{s: S_* \mid V(s) \leq H\}$. By the non-increasing constraint on p, d we have that for every s in S_H any transition decreases $V(s)$ by at least $d(v) \geq d(H) = d_H$ say, with probability at least $p(v) \geq p(H) = p_H$. Note that there does not have to be an actual s in S_* with $V(s) = H$ for this condition to apply.

Now fix s in S_H with therefore $V(s) \leq H$. The probability that V will eventually become 0 via transitions from that s is no less than $(p_H)^{\lceil H/d_H \rceil}$, since taking the probability-at-least- p_H option to decrease V by at least d_H , on every transition, suffices if that option is taken at least $\lceil H/d_H \rceil$ times in a row.

Since the above paragraph applies for all s in S_H , the probability of transitions’ escaping S_H eventually is bounded away from zero by $(p_H)^{\lceil H/d_H \rceil}$ uniformly for *all* of S_H . We can therefore appeal to the zero-one law [12],[23, Sec. 6],[20, Sec. 2.6], which reads informally

Let process P be defined over a (possibly infinite) state space S , and suppose that from every state in some subset S_H of S the probability of P ’s eventual escape from S_H is at least ε , for some fixed $\varepsilon > 0$. Then P ’s escape from S_H is *AC*: it occurs with probability one.

Note that the zero-one law applies even if S_H is infinite.

It is possible however that the escape occurs from S_H not by setting V to 0 but rather by setting V to some value greater than H , i.e. occurs “at the other end”. Because of possible nondeterminism, there might be many distributions describing the escape from S_H ; but because we know escape is *AC*, they will all be full distributions, i.e. summing to one. Let δ be any one of them.

Set $z = \delta_{S_0}$, i.e. so that the probability of indeed escaping to $V=0$ is z . Then the probability of escaping to $V > H$ instead is the complementary $1-z$ for that δ , and the expected value of V over δ is at least $z \times 0 + (1-z) \times H$, since the actual value of V in the latter case is at least H . But by *super-martingale*, we know that the expected value of V when escape occurs from S_H , having started from s , cannot be more than $V(s)$ itself. So we have $(1-z)H \leq V(s)$, whence $z \geq 1 - V(s)/H$.

Now we simply note that the inequality $z \geq 1 - V(s)/H$ holds for any choice of s, H and, in particular, having fixed our s we can make H arbitrarily large. [†] Thus z , the probability of escape to $V=0$, i.e. to S_0 , must be 1 for all s . ⁸ \square

⁸ A subtle issue here is that there might be $V=0$ states that s can reach via all of S_* but from which it is blocked because it must terminate when $V > H$ — and our z above does not take those into account. That is, the inequality wrt. z might apply only to a *subset* of the $V=0$ states that s can reach in the full system S_* . But the “actual z ”, i.e. for the full system, can only be greater still — and so the result holds regardless.

4.2 Discussion of the rule and the necessity of its conditions

What happens if V is not a super-martingale? Then *ACT* could be be (unsoundly) proved e.g. for a biased random walker biased away from 0, say $1/3$ probability of stepping closer to zero and $2/3$ of stepping away. Setting its variant equal to its distance from zero satisfies *progress*, but not *super-martingale*.

What happens without *progress*? Then a stationary walker would be compliant, satisfying *super-martingale* but not *progress*. (Remember our *super-martingale* does not require strict decrease: a stationary walker would satisfy it.)

Why not allow V to go below 0? In the proof we argued the expected value of V on exit from S_H would be at least $z \times 0 + h \times H$ — but it could be much lower if an exit in the zero direction could set V to a negative value.

In fact V can be boundedly negative: we would just shift the whole argument up. But V must be bounded below, otherwise the rule is unsound. Consider the “captured spline” example (in Fig. 7 of §7.7 below), and replace the 0-variants for escape by variants $-2(n+1)^2$. The ∇ rule (defined in §5) would now apply with ∇ the constant function -1 . For the current p, d rule we could use the large negative escape-variants to increase the (positive) along-the-spline variants so that they became unbounded.

What happens when V is bounded? Consider again the $2/3$ – $1/3$ biased random walk. We can synthesise a (super-)martingale by setting $V(n)=0$ when $n=0$ and solving for $V(n-1)/3 + 2V(n+1)/3 = V(n)$ otherwise — it gives the definition $V(n) = 2^n - 1/2^{n-1}$ with which *super-martingale* is satisfied by construction. Then, since V is injective, we can go on to define $p(v), d(v)$ to be the probability, decrease resp. actually realised by the process whenever its variant is v , appearing at first sight to satisfy *progress* trivially: set p to be the constant function $1/3$ and $d(v)$ to be $2-v$ in this example. Both p, d are non-increasing and strictly positive over variant values taken by the process.

But *progress* is not satisfied, because the functions d, p must be defined and non-increasing over *all* positive values v and, in particular, not only over variant values actually taken by the process: that is, they must be defined even for values v for which there is no s with $v=V(s)$. In this example $d(v)$ decreases to 0 as v approaches but never reaches 2, and so we cannot set a non-zero and non-increasing value for $d(2)$ itself. (In §7.7 a similar example is given where instead it is $p(2)$ that cannot be defined.)

The point in the proof at which this “any v whatsoever” is used is marked by a marginal (†), where we let H increase without bound. That H does not have to be $V(s)$ for any “actual” s .

In summary: if V is bounded but the values of “actual” d (or p) are not bounded away from zero, then for any H greater than all $V(s)$ there can be no non-zero value for $d(H)$ (or $p(H)$) and the proof fails.⁹

⁹ See Thm. 2 in §9.1 a for place where unbounded variant seems to be required.

Why are p, d functions of the variant rather than of the state?

Indeed they could have been defined as functions of the state (simply by composing them with V). In that case the non-increase conditions would become

If states s, s' are such that $V(s) < V(s')$ then $p(s') \leq p(s)$ and $d(s') < d(s)$.

But we would have to add that V over S_* must either take only finitely many values or be unbounded, because we would then no longer be considering the v 's that correspond to no s . That conflicts with our “purely local reasoning” goal.

Why not simply require V to be unbounded? For a finite state space V cannot be unbounded; yet for finite state spaces a termination argument is (usually) easy. As our rule stands, termination for finite state-spaces is handled as part of the general argument, not as a special case.

Are there alternatives formulations of *progress*?

Yes: there are several alternatives.

The rule (§2.2) uses *progress* in its proof (§4.1) only to show bounded-away-from-zero escape from an arbitrary but bounded V -region $(0, H]$ that we called S_H . That is, starting from any s in S_H the probability of reaching eventually an s' with either $V(s')=0$ or $V(s')>H$ is bounded away from 0, where the bound can depend on H . (It is *super-martingale* that then converts that to AC escape to S_0 alone, that is $V(s')=0$, by letting H increase without bound.) Any other condition with the same force would do, and a significant programming-oriented example is given in §5.

Another alternative, more suited to the situation where S, T are laid out as a transition system or as a Markov process (but not so suitable for systems expressed as programs), is simply to require that the V -image of S_* have no accumulation points. (An example of this kind of condition is found in [16, Item (i)] and [9, Condition (2) *proper divergence*].) In that case the size of the set $V(S_H)$, i.e. the “number of V 's” in any region $(0, H]$, is required to be finite for any H . If the system is deterministic (or at least only boundedly nondeterministic)¹⁰ then if in every transition V must decrease, by no matter how small an amount and with no matter how small a probability, escape from $(0, H]$ is assured because Zeno-effects cannot occur in a finite set: the $p(v), d(v)$ required in our rule (§2.2) can be synthesised by taking minima over the whole (finite) set $(0, v]$, i.e. with $H=v$.

Defining p, d everywhere, rather than only on “actual” v 's, is not a burden if the v 's are unbounded: define for example $\hat{p}(v')$ to be the infimum of $p(v)$ for all actual v 's with $v \leq v'$. Those extra values $\hat{p}(v')$ are never used, since there are no states with $V(s)=v'$: just the existence of the extra values is enough.

The only time this trick does not work is precisely the case we are discussing, where v is bounded but $p(v)$ tends to zero.

¹⁰ Both [16, 9] deal only with deterministic systems, i.e. *stationary* Markov processes.

5 An equivalent rule based on parametrised strict super-martingales

Pursuing the theme of equivalent formulations of *progress* (mentioned just above), we give here an equivalent rule in which *progress* is removed altogether, and replaced by parametrically strict *super-martingale* as follows:

There must be a non-increasing strictly positive function ∇ on the positive reals such that whenever we have $V(s)=v$ for some s, v and some δ (2) in $T(s)$ then $\mathcal{E}_\delta V \leq v - \nabla(v)$.¹¹

Call this formulation the “ ∇ rule”, and the original the “ p, d ” rule. Although the ∇ rule is simpler to state than the p, d rule, in practice the definition of ∇ can be complicated, often the definitions of p, d are more straightforward. The similarity of this rule with other strict super-martingale rules is clear: our condition is weaker (the rule stronger) because we do not impose a uniform ε across all of S_* .

We show first that the p, d rule implies this ∇ rule.

Lemma 1. (*Technical*) _____ Let f be a non-negative function over the non-negative reals, and let y, y' be non-negative reals; let δ be a discrete distribution on the non-negative reals. Then

$$\mathcal{E}_\delta f \leq y \quad \text{implies} \quad \delta_{\{x|f(x)<y'\}} \geq 1 - y/y'. \quad (3)$$

That is, if δ guarantees that the expected value of f is no more than some y , then for any y' we have that δ is guaranteed with probability at least $1 - y/y'$ to set f to a value no more than y' .¹²

Proof Let p be the aggregate probability that δ assigns to $\{x \mid f(x) \geq y'\}$. Then, since δ is fixed, the smallest possible value of $\mathcal{E}_\delta f$ is py' , found by making f itself as small as possible: that occurs when $f(x)=0$ for all x with $f(x) < y'$ and $f(x)=y'$ for all x with $f(x) \geq y'$. Thus $py' \leq \mathcal{E}_\delta f \leq y$, whence $p \leq y/y'$ and so the complementary $\delta_{\{x|f(x)<y'\}}$ is $1-p \geq 1 - y/y'$. \square

Lemma 2. *Guaranteed decrease of variant* Let V, S, T etc. be as above. Suppose for some state s in S_* we have that any T -transition is guaranteed to decrease the expected value of V by at least some $\varepsilon > 0$.

Then any T -transition is guaranteed with probability p to decrease the *actual* value of V by at least d , where $d := \varepsilon/2$ and $p := d/V(s) - d$.

Proof Let δ in $T(s)$ be a T -transition from s , and for Lem. 1 set $y = V(s) - \varepsilon$ and $y' = V(s) - \varepsilon/2$ and $f = V$. Then δ is guaranteed with probability at least

$$1 - \frac{y}{y'} = 1 - \frac{V(s) - \varepsilon}{V(s) - \varepsilon/2} = \frac{\varepsilon/2}{V(s) - \varepsilon/2}$$

¹¹ Note that ∇ must be defined on *all* the positive reals, not just on the variant values the process can actually take.

¹² If $y' \leq y$ then of course this guarantee is vacuous.

to decrease V by at least $V(s) - y' = \varepsilon/2$.

So we let d be $\varepsilon/2$ and p be $d/V(s) - d$. □

We can now conclude that the ∇ rule implies the p, d rule because if the ε in Lem. 2 is a non-increasing but never-zero function of $V(s)$, then the p, d -values synthesised there are also non-increasing never-zero functions of $V(s)$. Non-increase of d follows from the assumed non-increase of ε , and the non-increase of p follows from increase of V and non-increase of d .

For the opposite direction, that the ∇ rule implies the p, d rule, we again let V, S, T etc. be as above. we will replace variant V by $V' = f \circ V$ where f is a real-valued function that is

- non-decreasing,
- strictly concave and
- of non-increasing curvature.

That would be equivalently $f' \geq 0$ and $f'' < 0$ and $f''' \geq 0$, for which an example is logarithm.

Now for any state s and δ in $T(s)$, we know that with probability at least $\hat{p} = p(V(s))$ the δ -transition decreases $V(s)$ by at least $\hat{d} = d(V(s))$, and from *super-martingale* we know that that $\mathcal{E}_\delta V \leq V(s)$. Then because of the concavity of f , the smallest value of $V'(s) - \mathcal{E}_\delta V'$ occurs when δ sends exactly weight \hat{p} to (possibly several) s' with $V(s') = V(s) - d$ and exactly weight \hat{p}' to s'' with $V(s'') = V(s) + d'$ where \hat{p}' is $1 - \hat{p}$ and $\hat{p}\hat{d} + \hat{p}'\hat{d}' = 0$. (The construction of \hat{p}', \hat{d}' is to make \hat{d}' as big as possible while satisfying *super-martingale* wrt. V .)

Because of f 's concavity, that smallest value of $V'(s) - \mathcal{E}_\delta V'$ will be non-zero; and because the curvature is decreasing, and p, d are non-increasing functions of $V(s)$, it will be non-increasing wrt. increasing values of $V(s)$; because f is non-decreasing, that is equivalently non-increasing wrt. increasing values of V' .

6 Relation to the rule of Fioriti and Hermanns

Fioriti and Hermanns' rule [8] does not have our *progress* condition; instead they require uniform bounded-away-from-zero decrease of the expected value of the variant, that is with the same bound for the whole of S_* .

But in §5 we showed that our rule is equivalent to one without *progress*, i.e. where *super-martingale* has been strengthened to the ∇ rule at (2) above.

Fioriti and Hermanns' rule is then the special case of (2) where ∇ is the everywhere- ε constant function. Furthermore, since that rule is complete for systems with finite expected time to termination, the result above means our proposed rule is also complete for that class. But –as observed in §2.2– our rule also applies to the unbounded random walk, where the termination time is infinite.

For further discussions of completeness, see §9.

7 Examples of termination and non-termination

7.1 Symmetric unbounded random walk (terminates)

We mentioned in §2.2 that with variant the “distance from 0” and p, d the constant functions $1/2, 1$ respectively the *ACT* of the one-dimensional symmetric random walker is immediate. We also stressed our concern with source-level reasoning. Here we illustrate such reasoning for a random-walk program:

```

s := 1
while s ≠ 0 do
  s := s + 1 1/2 ⊕ s - 1
end

```

Reasoning in Kozen’s style [18] (here written in *pGCL* [25, 20]) would generate just these two elementary verification conditions for the proof-rule of §2.2:¹³

- The expected value of the variant does not increase:

$$\text{super-martingale} \quad s \geq \text{wp}.(s := s + 1_{1/2} \oplus s - 1).s$$
- With probability at least $1/2$ the variant decreases by at least 1:

$$\text{progress} \quad 1/2[s=N] \leq \text{wp}.(s := s + 1_{1/2} \oplus s - 1).[s \leq N-1]$$

The *wp* is the probabilistic generalisation of Dijkstra’s weakest precondition [6, 18, 25, 20].


To allay suspicions that might be raised by the simplicity of the above, we “unpack” the reasoning used in the proof of Thm. 1, showing in particular how the zero-one law contributes in this particular example. Without loss of generality we take the state-space to be the non-negative integers, start at position $s=1$ and show that eventually we will reach $s=0$.

Consider say the segment $1 \leq s \leq 100$ of the line, and the *bounded* random walk within it, beginning (as we said above) at $s=1$. Since s is decreased by $d=1$ with probability $p=1/2$ at every step, i.e. the *progress* property, the walker’s chance of moving to $s=0$ is at least $1/2^{100}$ for every $1 \leq s \leq 100$. Thus its escape from $[1, 100]$ is *AC*, whether that escape is high or low, and the expected value of s when that happens will be $z \times 0 + (1-z) \times 100$, that is $100(1-z)$, where z as before is the probability of escaping to $V=0$.

But the expected value of s is constant at 1 (the *super-martingale* property), no matter how many steps are taken, so that in fact $z=99/100$. That is, the probability that escape occurs to $s=0$ rather than to $s=101$ is $99/100$, establishing in any case that $s=0$ is reached from $s=1$ with at least that probability.

Now replay the argument within the segment $1 \leq s \leq 10^6$ instead. The walker’s behaviour is not affected by the segment within which we reason –it does not “know” we are looking at $[1, 10^6]$ – and it moves just as it did in the 100 case. But because we are thinking about 10^6 this time, our conclusions are strengthened to “escape from $s=1$ to $s=0$ with probability $1-1/10^6$ ”.


¹³ In fact they are both equalities; but in general the inequalities shown are what must be verified.



The p, d version of our rule (§2.2) establishes *ACT*. The variant is the distance from 0 which, everywhere except 0 itself, is a (super) martingale that decreases by at least $d=1$ with probability at least $p=1/2$.

Fig. 1. The unbounded symmetric random walk example

7.2 Constant-bias random walk (non-terminating)




Here the walker has constant bias away from 0, and indeed termination is not *AC*.

Although *super-martingale* is satisfied and we can define $p(v)=1/3$, it is impossible to define a non-increasing function d that gives a lower bound on the amount by which the variant decreases: the variant at State s is $2 - 1/2^{s-1}$, bounded above by 2 and forcing the non-increasing but strictly positive d impossibly to satisfy $d(2)=0$.

Fig. 2. The constant-bias random walk example

In Fig. 2 we have a one-dimensional random walk that does *not* terminate *AC*. If we synthesise a variant V that is an exact martingale, as shown, we satisfy *super-martingale* by construction. And its decrease occurs with probability (at least) $1/3$ everywhere. But because the variant is bounded, we cannot define a d that satisfies *progress*, so our termination rule does not apply. (And §5 shows that the ∇ -rule does not apply either.) In §§9.2,9.3,9.4 we see that in fact this walker does not terminate *AC*.



State s goes up with probability $(s+1)/(2s+1)$. It is strictly biased away from 0 everywhere.

Here we use the p, d version (§2.2) of the proof rule: the expected value of the variant after a transition is equal to its actual value before (except at State 0, where our rule does not require it to be). But still this walker is strictly biased away from 0 at all positions, with that bias however decreasing towards zero with increasing distance. In spite of that bias, still its termination is *AC*.

Fig. 3. The harmonic-bias random walk example

7.3 Harmonic-bias random walk (terminates)

In Fig. 3 we see a biased one-dimensional random walk that still terminates *AC*. The key point is that the bias decreases as distance from 0 increases, tending to “symmetric” in the limit.¹⁴

Here the variant is unbounded. (Compare §7.2 just above, where the variant is bounded.) Condition *super-martingale* is satisfied by construction; define $p(v) = 1/3$ everywhere; and define $d(v)$ to be $1/s$ where s is the largest such that Harmonic Number H_s is no more than v . This d is non-increasing in v and strictly positive.

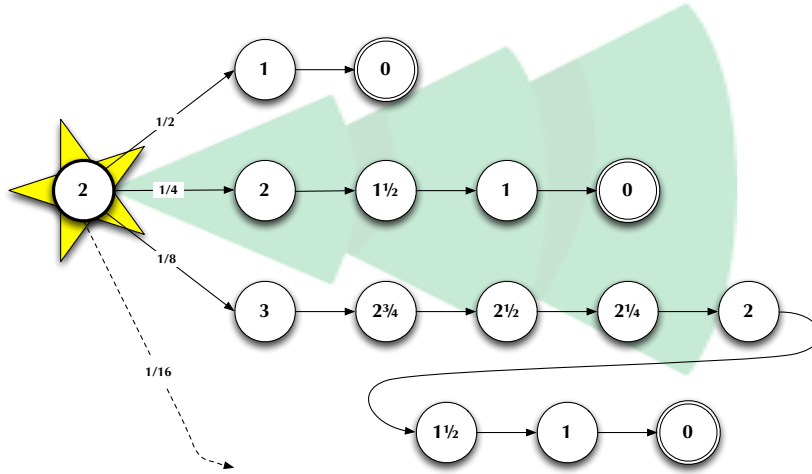
An alternative proof of termination for this process is provided by the general techniques of §9.1.

7.4 The “tinsel” process (terminates)

Here we exhibit a process whose infinite stopping time is obvious from its construction. (The random-walk process (§7.1) has the same infinitary property, but it is not so obvious.)

The root branches with probabilities $1/2, 1/4, \dots, 1/2^n, \dots$ to straight paths of length $2, 4, \dots, 2^n, \dots$ resp. each of whose contributions to the expected stopping time is therefore $(1/2^n)/(1/2^n) = 1$. Since there are infinitely many children of the root, the expected stopping-time overall is infinite. See Fig. 4, where the variant function for *ACT* is shown.

¹⁴ Although we constructed this example ourselves, we later found it in [9, Sec. 3(b)].



Variants in $(0, 1]$ decrease by at least 1 with probability at least 1. In $(1, 2]$ the (smaller) lower bound is $1/2$; in $(2, 3]$ it's $1/4$; in $(3, 4]$ it's $1/8 \dots$


The *super-martingale* condition is satisfied trivially except at the root node, where the small calculation $1/2^1 + 2/2^2 + 3/2^3 + \dots = 2 \leq 2$ is needed to see that it is satisfied there too.

The expected stopping time however is $\sum_{n \geq 1} 2^n / 2^n = \infty$. (We call it “tinsel” because it’s like long ribbons hanging down from a tree.)

Fig. 4. The “tinsel” process (rotated 90°)

7.5 The “curtain” process (terminates)

This variation on infinite stopping time begins with transitions that either move away from the root or “drop down” to ever longer straight runs. Again the stopping time is infinite but termination is still *AC*. See Fig. 5, where the variant function is shown.



Variants in $(0, 2]$ decrease by at least 1 with probability at least $1/2$. In $(2, 3]$ it's $1/3$; in $(3, 4]$ it's $1/7$. . . In $(s-1, s]$ it's at least $1/(2^{s-2}-1)$.

The *super-martingale* condition is satisfied trivially everywhere.


The expected stopping time however is again $\sum_{n \geq 1} 2^n / 2^n = \infty$, as for Fig. 4. (We call it “curtain” because many short runs hang down from a single long run.)

Fig. 5. The “curtain” process (again rotated 90°)

7.6 The escaping spline (terminates)

Here we illustrate in Fig. 6 how our rule depends on the actual transition probabilities in an intuitive way, that a “spline” whose overall probability of being followed forever is zero gives a variant with which we can prove its termination. (Complementarily, if the probability of remaining in the spline is not zero then our rule does not apply, as we show in §7.7.)

The states are numbered from $s=1$ at the left, and $V(s)$ is s itself. The function $p(v)$ is $1/v+1$ and the function $d(v)$ is 1 everywhere: at state $s \neq 0$ the variant is s , and with probability at least $1/s+1$ the value of the variant will decrease by at least 1. In fact, for most s with $V(s) \neq 0$ the variant by much more than 1 — the function d gives only a lower bound for the actual decrease in the variant.



Each horizontal transition has probability of one minus the (vertically downwards) escape immediately before it. Each variant $2, 3, 4, \dots$ turns out to be the previous variant divided by its incident probability, establishing *super-martingale* by construction. The successive probabilities of *not* having escaped are corresponding prefixes of the infinite product $1/2 \times 2/3 \times 3/4 \times \dots$ which tend to zero. Hence the variant increases without bound, proving that eventual escape is *AC*.

In general, if the product of the “stay on spline” probabilities tends to zero, the variants —the reciprocals of those prefix probabilities— increase without bound.

Fig. 6. The “escaping spline” process


7.7 The captured spline (non-terminating)

In the example of Fig. 7, based on [20, Sec. 2.9.1], the process does not escape with probability one. If we applied the strategy of the escaping spline (Fig. 7), we would choose variant $V(s) = 2s/s+1$. It is a (super-)martingale because in general

$$\begin{aligned}
 & V(s-1) \times 0 + (1 - 1/(s+1)^2) \times V(s+1) \\
 &= 1/(s+1)^2 \times 0 + (1 - 1/(s+1)^2) \times (2(s+1)/s+2) \\
 &= (s^2 + 2s/(s+1)^2) \times (2(s+1)/s+2) \\
 &= (s^2 + 2s/s+2) \times (2(s+1)/(s+1)^2) \\
 &= 2s/s+1 \\
 &= V(s) .
 \end{aligned}$$

The decrease function d is trivial: we can set it to the constant 1, since the potential decrease is always *at least* 1 with probability $1/(s+1)^2$.

But for $p(v)$ we choose $(2-v)^2/4$, i.e. a value that is no more than $1/(s+1)^2$ when $v = 2s/s+1$. Whatever that value is, it is clear that it approaches 0 as v approaches 2, and so we will not be able to select a non-zero value for $p(2)$. As for §7.2, the results of §§9.2,9.3,9.4 show that this process does not terminate *AC*.



As before, each horizontal transition has probability (this time not shown) of one minus the escape immediately before it; and each (speculative) variant is the previous variant divided by its incident probability. The successive probabilities of *not* having escaped are now corresponding prefixes of an infinite product $(1-1/2^2) \times (1-1/3^2) \times (1-1/4^2) \times \dots$ which, unlike the earlier one of Fig. 6, does not diverge: rather it converges to $1/2$. Hence eventual escape is with probability only $1-1/2 = 1/2$.

Making the variants the reciprocals of those cumulative escape probabilities, as in Fig. 6, results in increasing variants bounded above by 2, which does not satisfy *progress* for $p(v)$ when for example $v=2$.

In general, the strategy of Figs. 6,7 works just when the successive “not yet escaped” probabilities tend to zero, since that is exactly when the variants, their reciprocals, increase without bound.


Fig. 7. The “captured spline” process

7.8 The two-dimensional random walk (terminating but not proved)

In Fig. 8 we recall the one-dimensional random walk, but this time using a variant equal to the logarithm of (one plus) the walker’s distance from the origin and a ∇ -style *progress* condition. (Compare Fig. 1 in §7.1.) For better comparison with the two-dimensional version, we have made the walk unbounded in *both* directions. It suggests that the two-dimensional walker could be treated with the variant being based on the logarithm of the walker’s *Euclidean* distance from the origin. Again using the ∇ rule, we would have at least to show (something like) that for all integers x, y we have

$$\begin{aligned} & \log((x+1)^2+y^2) + \log((x-1)^2+y^2) + \log(x^2+(y+1)^2) + \log(x^2+(y-1)^2) \\ & < 4\log(x^2+y^2) . \end{aligned}$$

Unfortunately, numerical calculations show that this inequality fails near the $|x|=|y|$ lines. It seems that the log function bends too much, is “too concave”.



Here we use the ∇ -version (§5) of the proof rule: the expected value of the variant decreases by at least some fixed positive and non-increasing function of its current value: the expected decrease here is $\frac{1}{2} \log \frac{n^2}{n^2-1}$, a non-increasing function of $\log n$.

Fig. 8. The unbounded symmetric random walk example

We therefore “flatten things out a bit” by trying a double-log $\log(\log(\cdot))$ instead, a function still concave but less so, and we have indeed shown by similar numerical calculations that the corresponding inequality

$$\begin{aligned} & \lgg(x+1)^2+y^2) + \lgg((x-1)^2+y^2) + \lgg(x^2+(y+1)^2) + \lgg(x^2+(y-1)^2) \\ & < 4 \lgg(x^2+y^2) \end{aligned} \tag{4}$$

is satisfied for all integers x, y with $|x|, |y| \leq 10,000$.¹⁵

Our conjecture is that (4) holds for all integers x, y and, if it does, it would establish termination for the two-dimensional symmetric random walk using a single variant function.¹⁶

See §9.3 for evidence that there is a suitable variant function, even if it turns out not to be \lgg .

8 Compositionality

Following [8], by “compositionality” we mean the synthesis of an *ACT*-proof for a system that is composed of smaller systems for each of which we have an *ACT*-proof already. For now, we study this only briefly.

Suppose we have a “master” system M and a number of component systems $C_{1..N}$. System M has at least N termination states, at which its variant V_M is therefore zero; and each component system has a designated start state s_n where its variant function V_n takes some value v_n . The composite system is then made by “plugging in” each component system’s starting state s_n to some termination state of M .

The systems in Fig. 4 (Tinsel) and Fig. 5 (Curtain) are examples of this, except that for them the number of component systems is infinite.

In Tinsel, the master M is a single infinite branch leading with ever-decreasing probability $1/2^n$ to termination in exactly one step. Its component systems $C_{1..}$

¹⁵ We write \lgg for that function. Very close to the origin the it is undefined: but those cases can be adjusted manually.

¹⁶ As hoped, \lgg fails in the three-dimensional case.

are straight line processes each with stopping time $2^n - 1$. The overall stopping time of the combination is infinite.

In Curtain the master M is a straight-line system with a probability of $1/2$ of termination at each step; its expected stopping time is 2. The component systems C_1, \dots this time have termination times of $2^n - n$. Again the overall stopping time of the combination is infinite.

Although we did give termination proofs for these two systems, we cannot (i.e. at the moment we do not know how to) *synthesise* such proofs from the master's and the components' proofs when the number of components is infinite. But here is what we can do when the number of components is finite:

- \$ – Define v_C to be the maximum over all n of v_n , that is a number at least as great as the starting variant of any of the finitely many subsystems.
- Modify System M by adding v_C to its variant function V_H .
- Paste the starting node s_n of each System C_n into the appropriate termination state of M . (These are therefore no longer terminating states.)
- Set ∇_{\square} for the new, single system to be the pointwise minimum over all C_n and M of their individual ∇_n and ∇_M functions.

If the individual systems satisfied the ∇ rule with their separate ∇ functions, then the composite system will satisfy the rule with the single ∇_{\square} function. The use of finiteness was in two places:

- \$ – That the v_C added to V_H was finite. (It is a sup taken over all subsystems.)
- That the ∇_{\square} is nowhere zero.

In Tinsel (Fig. 4) and Curtain (Fig. 5), having infinitely many components, the failure of synthesis occurs at the two points \$, because v_C is infinite. In spite of that, as the examples show, we were able to find proofs “by hand” (i.e. not synthesised). Note however if a proof method were complete only for finite stopping-time systems, there can be no synthesis in these two cases: although all the component systems have finite stopping times, but the composite systems do not.

9 Related historical results on Markov chains

9.1 The work of Blackwell: random walks and radially symmetric trees

Blackwell [1] gives a general technique for proving termination of a certain subclass of Markov processes, those moving both down *and up* so-called “radially symmetric trees”. (It also provides an independent proof of termination for our example §7.3, the harmonic random walk.)

Definition 1. *Radially symmetric tree* _____ A radially symmetric tree is finitely branching, having the property that each node at depth d has exactly c_d children, where the root has depth zero and all the c_0, c_1, \dots are integers at least 1. □

A radially symmetric tree is infinite, and has no leaves.

Definition 2. *Random tree-walk* _____ A random walk on a radially symmetric tree starts at any node and chooses uniformly to move either to its parent or to one of its children, thus with probability $1/c_d+1$ for a node at any positive depth d along any of its connecting arcs. At the root, where there is no parent, the probability is instead $1/c_0$. (Only the root has no parent.) Termination occurs when the root is reached. \square

Radially symmetric trees are determined uniquely by their c_d 's, and examples include the following:

- (i) Each node has exactly one child, thus a single path starting at the root.
- (ii) Each node has exactly two children, thus an infinite binary tree.
- (iii) Each node has either one or two children, according to the scheme that c_d is two just when d is a power of two (and thus is one otherwise).

The first example generates the one-sided symmetric random walk in one dimension, and it terminates with probability one (§7.1). The second example is (effectively) a collection of $2/3-1/3$ asymmetric random walks in one dimension, down the branches of the tree: its termination is not *AC* (§4.2). The third, more complicated example is a binary tree with only infrequent splittings: we show in Fig. 9 below that it terminates, and quote a completeness result for such trees.

The Blackwell-style proof of (iii)'s termination (a sanity check of our claim) follows [19], and is ultimately based on Blackwell's Thm. 2 below. Note that it is an if-and-only-if:

Theorem 2. *Blackwell [1]* _____ Let p_n for $n \geq 0$ be a sequence of probabilities, with $p_0=1$, and consider the Markov matrix on the non-negative integers defined $M_{n,n+1} = p_n$ and $M_{n,n-1} = q_n = 1-p_n$.

Then any Markov chain with this matrix will eventually reach 0 almost surely *if and only if* the equation $f(n) = q_n f(n-1) + p_n f(n+1)$ has no non-constant bounded solution. ¹⁷ \square

A corollary of Thm. 2 gives us a classification of *ACT* for radially symmetric trees:

Corollary 1. *Lessa [19]* _____ Given a radially symmetric tree, define variant function

$$V(d) := \sum_{0 \leq i < d} \frac{1}{c_0 \times \cdots \times c_i} \quad (5)$$

where d is a depth of some node in the tree and c_d is the number of children for nodes at that depth. Every node at depth d has variant $V(d)$, and this V is indeed a super-martingale.

The random walk on this tree, defined as in Def. 2, terminates everywhere if and only if $V(d)$ is unbounded as $d \rightarrow \infty$.

Proof \square

¹⁷ Note however that our rule does not require V to be unbounded. See §4.2.

Lessa's proof of Cor. 1 uses Blackwell's theorem Thm. 2; but our variant rule here provides an independent proof for Fig. 9, i.e. without Blackwell's theorem. More significantly however, Blackwell's theorem provides us with a completeness result for using our variant rule, at least for one-dimensional random walks.

9.2 Blackwell's completeness result

Blackwell's work [1] on classifying recurrence in Markov processes suggests how we might understand the coverage of our new rule. He considers Markov processes with countable state spaces and stationary (i.e. fixed) transition probabilities, and shows that such processes have essentially unique structures of recurrent and transient sets. We now give a summary, using (partly) Blackwell's terminology as well as what we have used elsewhere in the paper.

Let C be a subset of the state space, and fix some initial state \hat{s} . Say that C is *almost closed* (wrt. that \hat{s}) iff the following conditions hold:

1. The probability that C is entered infinitely often, as the process takes transitions starting from \hat{s} , is strictly greater than zero and
2. If C is visited indeed visited infinitely often, starting from \hat{s} , then eventually it remains within C permanently.

Say further that a set C is *atomic* iff C does not contain two disjoint almost-closed subsets.

Finally, call a Markov process *simple atomic* if it has a single almost-closed atomic set such that once started from \hat{s} it eventually with probability one is trapped in that set. We then have


Theorem 3. *Corollary of Blackwell's Thm. 2 on p656) [1]*

A Markov process is simple atomic (as above) just when the only bounded solution of the equation $\mathcal{E}_\delta V = V(s)$, that is Blackwell's Equation (his 6), stating (in our notation) that V is exact, neither super- nor sub, is constant for all s in S_* and δ in $T(s)$. \square

We adapt the above to our current situation as follows. As above, we fix a starting state \hat{s} , and we collapse our termination set S_0 to a single state s_0 , adjusting T accordingly and in addition making T take s_0 to itself. We then assume that the probability of \hat{s} 's reaching s_0 is one. We now note:

- (1) Our termination set $\{s_0\}$ is almost-closed and atomic, because
 - (i) almost closed: Our process reaches s_0 with non-zero probability (in fact we assumed with probability one) and, once at s_0 , it remains there.
 - (ii) atomic: Our set $\{s_0\}$ has no non-empty subsets.
- (2) We now recall that in fact s_0 is reached with probability one, so that the whole process is simple atomic.
- (3) From our Thm. 3 we conclude that the only possible non-trivial variant is unbounded.

Thus –in summary– we have specialised Blackwell's result to show that if we have a non-trivial exact variant that is bounded, then the process *does not* terminate AC. This is a result in the style of Chatterjee et al. [5]. (See §2.3.)



Every node at depth 2^n has two children; all others have one child. Transitions from a node are uniformly distributed over its arcs, thus $1/3$ for each of two children and $1/3$ up, and otherwise $1/2$ up and $1/2$ down.

The variant function generated according to the scheme of (5) is

depth d	0	1	2	3	4	5	6	7	8
c_d	1	2	2	1	2	1	1	1	2
$V(d)$ from (5)	0	$1/1$	1	$3/2$	$7/4$	2	$17/8$	$9/4$	$19/8$
			$+ 1/1 \times 2$	$+ 1/2 \times 2$	$+ 1/4 \times 1$	$+ 1/4 \times 2$	$+ 1/8 \times 1$	$+ 1/8 \times 1$	$+ 1/8 \times 1$
simplified	0	1	$3/2$	$7/4$	2	$17/8$	$9/4$	$19/8$	$5/2$

At Depth 4, for example, we have $1/3 \times 7/4 + 2/3 \times 17/8 = 2$, thus satisfying *supermartingale* at that position. Because the variant at depth 2^d is $1 + d/2$, i.e. increases without bound, it is straightforward to construct functions p, d , showing that this process terminates.

Fig. 9. Blackwell's radially symmetric tree

9.3 The work of Foster: completeness

Foster [9] gives a characterisation of Markov processes for which a technique like ours is guaranteed to work. A significant example is the two-dimensional symmetric random walk, supporting our conjecture in §7.8.

Because Foster’s paper seems quite technical, we give here a “translation” into our own terms. His equations will be referred to as (F1) etc. and his sections as §F1 etc.

— **§F1** We assume that our state space S is countable, enumerated $s_0, 1, \dots$ with the termination subset S_0 being just a single point $\{s_0\}$, and we extend our transition function T to all of S , i.e. not just over S_* , by making it take s_0 to itself. The enumeration should correspond roughly to “being further from s_0 ”, which is made precise in conditions (F6) and (F8).

Foster is concerned with conditions for the existence of a super-martingale (F1) variant function V from S to the non-negative reals (F3), where V is unbounded without accumulation points (F2).

We assume that T is deterministic, and thus specialise it to be of type S to $\mathbb{D}S$ rather than to $\mathbb{P}S$.

— **§F2** The “limit” of taking transitions forever is defined to be T^* , say, using the “Cesaro average”¹⁸ that avoids the problem of recurrence when considering simply T^0, T^1 etc. composed in the Markov style.

But (F4) is not as simple as it looks. It seems to imply that there is no infinite chain of transient states (such as in the “spline” examples). For if there were, the mass travelling down the chain would be “lost” in the Cesaro average, and the sum would not be one. This turns out to be important in the discussion of (4’) below.

Kendall’s [16] earlier result is

If there is a variant as in §9.3, then T^* takes any starting state to a full distribution on S (i.e. not partial), and there is a *finite* subset C of S from which T does not escape.

Foster then explains that the current paper’s purpose is to explore the opposite implication to Kendall’s [16], i.e. that, under “certain weak additional assumptions” on T , if there is such a (finite) subset C as above then there is a V satisfying the conditions of §9.3.

His additional assumptions include (by implication) that C is reached with probability one from anywhere in S , his (F4’), because he argues that (F4, F5, F6) together imply (F4’). That looks at first like the zero-one law. But note that (F6) does not bound the probability of escape away from zero: it merely says that it is not zero, and that is not usually enough. Together with (F4) though it suffices,

¹⁸ <http://www.sciencedirect.com/science/article/pii/0304414977900321> .

because (F4) seems to say any transient state (even if there are infinitely many) must be visited infinitely often (since otherwise the mass moving among the transient states would be “lost”, as suggested above).

His additional assumptions are then

- (F6) From any state in S_* there is a non-zero probability of reaching $S_0=\{s_0\}$ eventually.
 - (F7) From any state in s_i in S_* there is for any $j>i$ a non-zero probability of reaching any s_j eventually. Note that s_j is in S_* also.
 - (F8) There is a single probability $\delta<1$ for the whole system such that for any N there is an i such that for all $j\geq i$ the state s_j cannot reach C within N steps and with probability at least δ . As he says, it’s a “remoteness” condition, intuitively mandating that the higher the i the longer it takes s_i to reach C with some fixed-beforehand probability δ .
- He notes that because of (F4′) the N (which depends on i) is finite: from s_i you must get to C eventually with probability at least δ because, in fact, you get there with probability 1.

— Statement of Theorem F2, and its application to the two-dimensional symmetric random walk

Recall that S is assumed to be countably infinite. Foster’s Theorem F2 reads

If T satisfies conditions (F4–8), then there is a variant function V on S that satisfies (F1,F3), i.e. that it is a non-negative super-martingale and (F2) that it tends to infinity as the state-index tends to infinity.

We note that condition (F2) implies that V is without accumulation points.

The implications of this theorem seem to be e.g. that there must be a variant in our style for the two-dimensional symmetric random walk, even if it has not (yet, as far as we know) been given in closed form. We check the conditions one-by-one:

- (F4) This is (apparently) replaced by (F4′).
- (F4′) The probability of reaching S_0 , that is the origin, is one everywhere.
- (F5) Once you are at the origin, you do not leave.
- (F6) Every state in S_* can reach S_0 with non-zero probability.
- (F7) Here we need an enumeration of the states: Foster uses the Manhattan distance, which makes nested “diamonds” . But since *every* state in S_* can reach every other state, in fact we do not need the enumeration yet.¹⁹
- (F8) Any state at Manhattan distance N cannot reach the origin at all until Step N , and so $\delta=0$ should do for this provided we assign higher indices to higher-Manhattan-distance states, which is what Foster does in §F3.

¹⁹ Foster’s (weaker) condition requires only that each state can reach every *higher-enumerated* state.

Applying Theorem F2 then gives us a super-martingale V such that $V(s_i)$ tends to infinity as i tends to infinity, which means that $V(s)$ tends to infinity as s gets Manhattan-further from the origin, given the indexing that we (i.e. Foster) have chosen.

To show that our rule applies, we need however to establish a progress condition. (See our earlier remarks about alternative progress conditions, in §4.2.) First define $p(v)$ to be $1/4$ for all v . Then for d , first consider the subset $S_{\leq v}$ of S comprising all those s with $V(s) \leq v$. Because of (F2) the V -image $V(S_{\leq v})$ of $S_{\leq v}$ must be finite; so set $d(v)$ to be the minimum non-zero distance between any two of them, that is $(\min V(s') - V(s) \mid s, s' \in V(S_{\leq v}) \wedge V(s') > V(s))$.

Thus there is guaranteed to be a V satisfying *super-martingale* and *progress* that establishes termination for the two-dimensional symmetric random walk (§7.8) — even if we don't know what it is in closed form. Foster's general proof is by construction, and we sketch it in App. A.

— Why Theorem F2 does not synthesise a variant for the three-dimensional symmetric random walk

Foster remarks [9, p. 590] that synthesis cannot succeed for the three-dimensional random walk (since it is known that it is not *ACT*); but he does not say which of his Theorem F2's conditions are not satisfied.

Clearly his (F4') is not satisfied (that the process is *ACT*); but that is a derived condition, a consequence of his original (F4–8), and so it is fair to ask which of those original conditions fails in three dimensions. Furthermore, his synthesis procedure is well defined whether the process satisfies *ACT* or not, and so we can therefore ask as well what is wrong with the V it synthesises for the three-dimensional random walk, in our terms.

For the first, it is condition (F4) that must fail. The process satisfies (F5), that the process is trapped at the origin, and (F6), that the origin is accessible from every point, and (F7), that every point is accessible from every other (except the origin). And it seems likely that it satisfies (F8), since by numbering the states in rings it can be arranged that higher-numbered states have arbitrarily high first-arrival times at the origin.

Failure of (F4) means that there is some bounded away from zero probabilistic mass that follows an infinite (not looping) path through the state space: that is the only way in which the Cesaro limit can “lose mass”, making (in Foster's notation) the sum $\sum_{j=0}^{\infty} \pi_{ij}$ strictly less than one.

For the second, the problem with the synthesised V is that it is bounded, a failure of condition (F2). Item 4. in §9.4 below shows that in that case our condition *progress* cannot be satisfied for that V .

9.4 A modern alternative to Blackwell's completeness argument

The result of §9.2 can be obtained much more directly using the program semantics of this paper, and an argument in the style of Thm. 1.

In [20, Lem. 7.3.1] we show (using the terminology here) that if V is a sub-martingale and is bounded,²⁰ ²¹ then if escape from S_* is *AC* (i.e. the loop terminates with probability one), the expected value of V on S_0 (i.e. on the states where the loop-guard is false) is at least the actual value of V in the initial state (of the loop).²²

Now if the initial state is in S_* then the value of V there is strictly positive; yet if escape occurs with probability one, the expected value of V on termination will be 0, since it is confined to S_0 . That contradiction means that we cannot have sub-martingale V be bounded and still terminate almost surely.

Here are some remarks on the relation between our argument and Blackwell’s Thm. 3. Blackwell states that a process is *ACT* just when its only bounded exact martingale is constant; our result just above states that an *ACT* process cannot have a bounded sub-martingale.

1. *What role does the “or is constant” criterion, from Blackwell’s theorem, play in our argument?*

Because we require V to be zero on S_0 , a constant V for us would be zero on all of S , meaning that S_* was empty (since V must be strictly positive there). So we should add to our result “unless S_* is empty.”

2. *Where do we use in our argument that V is bounded?*

How does our argument fail if we don’t?

We use it in our appeal to [20, Lem. 7.3.1], where in fact V is assumed to be between 0 and 1. If V is simply bounded above (but not by one, necessarily), then it can easily be brought within range by scaling. If V is unbounded, however, it cannot be brought within range that way.

An easy counter-example is the symmetric random walk starting at 1 and aiming to reach 0. The variant “distance from 0” is an exact martingale, and has value 1 on initialisation. But it is unbounded, and so the conclusion “its (expected) value on termination is at least its starting value” is false. Indeed on termination its (actual) value has decreased from 1 to 0.

3. *What’s an intuitive (and easy to understand now, in retrospect) reason that our conclusion is “obvious”, without appealing either to Blackwell or to [20]?*

Think of the variant value V , initially concentrated at s , as being gradually “dissipated” whenever some probabilistic weight escapes to S_0 . Since V is zero there, the sub-martingale property requires that V increase, to compensate, on the remaining probabilistic weight still within S_* . But because V is bounded, that increase cannot go far enough — it eventually must stop. And that means that some probabilistic weight remains trapped within S_* .

²⁰ Note we say *sub*-martingale, i.e. that the expected value of V can increase.

²¹ In that part of [20] we are working an invariant, here our V , that is bounded by 1. That loses no generality, since any bounded V can be divided by its least upper bound without disturbing its sub-martingale property.

²² See §B for a précis of this loop rule.

4. Our rule §2.2 proves *ACT* given super-martingale and progress for some V . Yet above we show that if V is a sub-martingale and bounded, then *ACT* cannot hold. Does that mean that progress cannot hold for any bounded, non-zero exact martingale?

Yes, it does mean that. If V is bounded, then when the process is at (or sufficiently near) V 's upper bound either

- (a) The function d must be arbitrarily small (tending to 0), and so it cannot be non-increasing with respect to a non-zero d -value taken above V 's upper bound (for which see Fig. 2), or
- (b) The function d is bounded away from zero, in which case the expected value of V must strictly decrease, thus not realising an *exact* martingale.

9.5 General comparison with refutation methods

Blackwell's result Thm. 3 says that a Markov process is atomic and simple if and only if all exact martingales are constant or unbounded. We showed (in our terms) that when a program terminates with probability 1, the termination set implies the program is atomic and simple (as a Markov Process). Then, using Blackwell's, result we are able to conclude that all exact martingales are constant or unbounded. In an independent proof (the one-liner §9.4) we can show this directly without going through Blackwell, namely that if a program has a non-constant exact martingale then it can't terminate with probability 1.

Chatterjee et al. also look at repelling super-martingales to refute almost termination. Their *Theorem 6* uses an ε -repulsing super-martingale with $\varepsilon > 0$ to refute almost sure termination. Their *Theorem 7* uses an ε repulsing super-martingale with $\varepsilon \geq 0$ to refute finite expected time to termination: i.e. to refute finite expected time to termination only a martingale is required.

Our result in §9.4 implies a new refutation certificate for programs: if the martingale is finite and non-constant it actually refutes termination with probability 1, not just finite expected time to termination.

For example, if we consider the one-dimensional random walker §7.1 it has an exact *unbounded* martingale, and therefore our rule shows that it terminates with probability 1. The walker in §7.2 has an exact *bounded* martingale, and this we can conclude does not terminate with probability 1. In both cases Chatterjee's Theorem 7 would deduce that neither have finite expected time to terminate.

10 Conclusion

Our overall aim is (has always been) to allow and promote rigorous reasoning at the level of program source-code. In this paper we have proposed a new rule, combining earlier ideas of our own with important innovations of others, and have attempted to formulate it in a way that indeed is will turn out to be suitable for source code.

That is, we hope that as an extension of what's here we will be able to formulate these rules in the program logic *pGCL*, or similar; and if the techniques are further extended subsequently, we would hope to do the same for those too.

Program logic also provides a rigorous setting not only for use of the rules but also for their proof in the first place. Although we did not use program logic here, for our proofs, we believe it would be possible e.g. in the style of [20].

Finally, we have left an intriguing open question: is there an elementary variant for the two-dimensional random walk? Foster [9] shows that there is such a variant, but he does not give it in closed form. We conjecture that lgg suffices, but have only verified that experimentally. Will we ever be able to set as a student exercise

Assuming the properties [...] of the function [...], use probabilistic assertions in the source code of the following program to prove that it terminates with probability one for any initial integers X, Y :

```

x,y:= X,Y
while x≠0 ∧ y≠0 do
  x,y:=    x+1,y
          ⊕ x-1,y
          ⊕ x,y+1
          ⊕ x,y-1
end ,

```

where iterated \oplus is shorthand for uniform choice (in this case $1/4$ each).

Acknowledgements

We are grateful to David Basin and the Information Security Group at ETH Zürich for hosting the six-month stay in Switzerland, during part of which we did this work. And thanks particularly to Andreas Lochbihler, who shared with us the probabilistic termination problem that led to it.

References

1. David Blackwell. On transient Markov processes with a countable number of states and stationary transition probabilities. *Ann. Math. Statist.*, 26:654–658, 1955.
2. Orieta Celiku and Annabelle McIver. Compositional specification and analysis of cost-based properties in probabilistic programs. In *Proceedings of Formal Methods*, 2005.
3. Aleksandar Chakarov. *Deductive Verification of Infinite-State Stochastic Systems using Martingales*. PhD thesis, University of Colorado at Boulder, 2016.
4. Aleksandar Chakarov and Sriram Sankaranarayanan. Probabilistic program analysis with martingales. In *International Conference on Computer Aided Verification*, 2013.
5. Krishnendu Chatterjee, Petr Novotný, and Dorde Žikelić. Stochastic invariants for probabilistic termination. arXiv:1611.01063.

6. E.W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
7. Javier Esparza, Andreas Gaiser, and Stefan Kiefer. Proving termination of probabilistic programs using patterns. In *International Conference on Computer Aided Verification*, 2012.
8. Luis Fioriti and Holger Hermanns. Probabilistic termination: Soundness, completeness, and compositionality. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2015.
9. F. G. Foster. On markov chains with an enumerable infinity of states. *Mathematical Proceedings of the Cambridge Philosophical Society*, 48(4):587–591, Oct 1952.
10. Friedrich Gretz, Joost-Pieter Katoen, and Annabelle McIver. Operational versus weakest pre-expectation semantics for the probabilistic guarded command language. *Perform. Eval.*, 73:110–132, 2014.
11. Probabilistic Systems Group. Collected publications.
www.cse.unsw.edu.au/~carrollm/probs.
12. S. Hart, M. Sharir, and A. Pnueli. Termination of probabilistic concurrent programs. *ACM Trans Prog Lang Sys*, 5:356–80, 1983.
13. Jifeng He, Karen Seidel, and AK McIver. Probabilistic models for the guarded command language. *Science of Computer Programming*, 28:171–92, 1997. First presented at FMTA '95, Warsaw.
14. C. Jones and G. Plotkin. A probabilistic powerdomain of evaluations. In *Proceedings of the IEEE 4th Annual Symposium on Logic in Computer Science*, pages 186–95, Los Alamitos, Calif., 1989. Computer Society Press.
15. Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. Weakest precondition reasoning for expected run-times of probabilistic programs. In *Proceedings of ESOP*, 2016.
16. David G. Kendall. On non-dissipative markoff chains with an enumerable infinity of states. *Mathematical Proceedings of the Cambridge Philosophical Society*, 47(3):633–634, 007 1951.
17. Konrad Knopp. *Theory and Application of Infinite Series*. London, 1928.
18. D. Kozen. A probabilistic PDL. *Jnl Comp Sys Sci*, 30(2):162–78, 1985.
19. Pablo Lessa. Recurrence vs transience: An introduction to random walks.
20. A.K. McIver and C.C. Morgan. *Abstraction, Refinement and Proof for Probabilistic Systems*. Tech Mono Comp Sci. Springer, New York, 2005.
21. A.K. McIver, C.C. Morgan, and T.S. Hoang. Probabilistic termination in *B*. In D. Bert, J.P. Bowen, S. King, and M. Walden, editors, *Proc ZB '03*, volume 2651 of *LNCS*, pages 2–6–239. Springer, 2003.
22. Annabelle McIver and Carroll Morgan. A new rule for almost-certain termination of probabilistic- and demonic programs. <https://arxiv.org/abs/1612.01091v1>, December 2016.
23. C.C. Morgan. Proof rules for probabilistic loops. In He Jifeng, John Cooke, and Peter Wallis, editors, *Proc BCS-FACS 7th Refinement Workshop*, Workshops in Computing. Springer, July 1996.
ewic.bcs.org/conferences/1996/refinement/papers/paper10.htm.
24. C.C. Morgan and A.K. McIver. Almost-certain eventualities and abstract probabilities in the quantitative temporal logic *qTL*. *Theo Comp Sci*, 293(3):507–34, 2003. Available at [11, key PROB-1]; earlier version appeared in CATS '01.
25. C.C. Morgan, A.K. McIver, and K. Seidel. Probabilistic predicate transformers. *ACM Trans Prog Lang Sys*, 18(3):325–53, May 1996.
doi.acm.org/10.1145/229542.229547.
26. Roberto Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.

A — Sketch of Foster’s proof [9]

We use the notation and definitions from §9.3 to present Foster’s Theorem 2.

Recall that we have assumed that $S_0 = \{s_0\}$, i.e. that termination occurs in a single state, and that we have adjusted (the assumed deterministic) T so that it takes s_0 to itself.

Write $f_i^{(t)}$ for the probability that T started from s_i reaches s_0 for the first time in the t -th step and (as Foster does) write p_{ij} for $T(i)(j)$, the probability that T takes one step from s_i to s_j ; more generally write $p_{ij}^{(t)}$ for the probability that T takes t steps to do that. Foster remarks just after (F9) that a simple special case is where time-to-termination is bounded, but notes that such an assumption excludes the symmetric random walk and moves immediately to the more general case.²³

For the more general case we note first that for $i > 0$ we have $f_i^{(t+1)} = \sum_j p_{ij} f_j^{(t)}$. So if we were hopefully to proceed simply by setting $V(s_0) = 0$ and $V(s_i) = \sum_{1 \leq t} f_i^{(t)}$ for $i > 0$, then in the latter case we would check the supermartingale property (F1) by calculating

$$\begin{aligned}
 & \sum_j p_{ij} V(s_j) \\
 = & \sum_j p_{ij} \sum_{1 \leq t} f_j^{(t)} \\
 = & \sum_{1 \leq t} \sum_j p_{ij} f_j^{(t)} \\
 = & \sum_{1 \leq t} f_i^{(t+1)} && \text{“above and } i > 0\text{”} \\
 \leq & \sum_{1 \leq t} f_i^{(t)}, && \text{“(actually equal unless } f_i^{(1)} > 0\text{)”} \\
 = & V(s_i),
 \end{aligned}$$

so that V would in fact be an exact martingale in most of S_* .²⁴ But this looks too good to be true, and indeed it is: in fact $\sum_{1 \leq t} f_i^{(t)} = 1$ by assumption, so this is just the special case where V is 1 everywhere except at s_0 ; and the martingale property is exact everywhere, except at states one step away from s_0 . And this trivial V does not satisfy *progress*.²⁵

Still, the above is the seed of a good idea. Using “a theorem of Dini” [17, Foster’s citation (4)],²⁶ that

If c_n is a sequence of positive terms with $\sum_n c_n < \infty$, then also

$$\sum_n \frac{c_n}{(c_n + c_{n+1} + \dots)^\alpha} < \infty$$

when $\alpha < 1$,

²³ [8] also treats the bounded-termination case explicitly.

²⁴ Think of the symmetric random walk, where everywhere-1 is an exact martingale except when $|x|=1$, where it is a proper super-martingale.

²⁵ It is trivial in Blackwell’s sense [1], a constant solution.

²⁶ There seems to be a typographical error here in Foster’s paper, where he writes $\sum_{r=1}^\infty \lambda^{(r)} f_i^{(r)}$ instead of $\sum_{r=1}^\infty \lambda^{(r)} f_1^{(r)}$.

Foster *increases* the $f_i^{(t)}$ terms above by dividing them by $\sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots}$, which is non-zero but no more than one,²⁷ and still (as we will see) the new, larger terms still have a finite sum. (A minor detail is that he must show that the sum $f_1^{(t)} + f_1^{(t+1)} + \dots$ does not become zero at some large t and make terms from then on infinite: his assumption (F7) prevents that by ensuring that from no state in S_* does a single T -step go entirely into S_0 .) With the revised V replacing the earlier “hopeful” definition, the calculation above becomes instead

$$\begin{aligned}
& \sum_j p_{ij} V(s_j) \\
= & \sum_{j \geq 1} p_{ij} \sum_{1 \leq t} f_j^{(t)} / \sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots} && \text{“revised definition of } V, \\
& && \text{and } V(s_0)=0\text{”} \\
= & \sum_{1 \leq t} \sum_j p_{ij} f_j^{(t)} / \sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots} \\
= & \sum_{1 \leq t} f_j^{(t+1)} / \sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots} \\
= & \sum_{1 \leq t} f_j^{(t+1)} / \sqrt{f_1^{(t+1)} + f_1^{(t+2)} + \dots} && \text{“denominator is not increased”} \\
\leq & \sum_{1 \leq t} f_j^{(t)} / \sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots} \\
= & V(s_i) .
\end{aligned}$$

This is encouraging: but we still must prove (F3) for our revised definition²⁸

$$V(s_i) = \sum_{1 \leq t} \frac{f_i^{(t)}}{\sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots}}, \quad \text{for } i \geq 1 \quad (6)$$

i.e. that it’s finite for all i and not only for the $i=1$ that Dini gave us; and we must show that it satisfies (F2), i.e. that it tends to infinity as i does.

For the first, Foster proves that $V(s_i) \leq V(s_1)/p_{1i}^{(t')}$ for any $i > 1$ and some $t' > 0$ with $p_{1i}^{(t')} > 0$, which is one place he uses (F7), in particular that every s_i is accessible from s_1 .

Specifically, he reasons as follows:

- (i) For that t' and any t we have $f_1^{(t'+t)} \geq p_{1i}^{(t')} f_i^{(t)}$, because we know that s_1 ’s journey to s_0 can go via s_i .
- (ii) The numerator $f_i^{(t)}$ in (6) can therefore be replaced by $f_1^{(t'+t)}/p_{1i}^{(t')}$ provided (\leq) replaces the equality.
- (iii) The sum in the denominator of (6) can be adjusted to start at $t'+t$ rather than t , still preserving the inequality.
- (iv) The overall sum in (6) of non-negative terms for $V(s_i)$ is now the “drop the first t' terms suffix” of that same sum for $V(s_1)$, which we already know to be finite (from Dini), but divided by $p_{1i}^{(t')}$ which we know to be non-zero.

For the second, Foster uses the δ from (F8), showing that $V(s_i)$ is at least $(1-\delta)/\sqrt{f_1^{(t_i)} + f_1^{(t_i+1)} + \dots}$ where t_i is the number of steps after which s_i reaches s_0

²⁷ It is the square-root of the probability that s_1 does not reach s_0 in fewer than t steps.

²⁸ Note the f ’s in the denominator are subscripted “1”, not “ i ”.

with probability at least δ for the first time. By (F8) that t_i tends to infinity as i does, and thus so does $V(s_i)$.

His detailed reasoning is as follows:

- (i) Since t_i 's tending to infinity is all that is required, any at-most-finite number of i 's where $t_i=0$ can be ignored. Thus pick $t_i \geq 1$.
- (ii) Then $V(s_i)$ is at least $\sum_{1+t_i \leq t} f_i^{(t)} / \sqrt{f_1^{(t)} + f_1^{(t+1)} + \dots}$, a suffix of its defining series (6).
- (iii) Since the denominators only decrease, we can replace all of the denominators by $\sqrt{f_1^{(t_i)} + f_1^{(t_i+1)} + \dots}$ while making the sum only smaller.
- (iv) From (F8) however and the choice of t_i we know that $\sum_{t_i \leq t} f_0^{(t)}$ is no more than $1-\delta$. Thus similarly we can replace $f_i^{(t)}$ by $1-\delta$ and remove the summation.
- (v) We are left with $V(s_i) \geq (1-\delta) / \sqrt{f_1^{(t_i)} + f_1^{(t_i+1)} + \dots}$, as appealed to above.

That completes the proof sketch.

B Loop rule précis (from §9.4)

To be clear, we interpret here our Lem. 7.3.1 from [20], which includes demonic nondeterminism. The lemma reads

Let invariant I satisfy $[G]*I \Rightarrow \text{wlp.body}.I$. Then $I \& T \Rightarrow \text{wp.loop}([\overline{G}]*I)$, where T is the termination probability $\text{wp.loop}[\text{true}]$ of the loop.

We note that

- G is a predicate over the variables of the program.
- Square brackets $[\cdot]$ convert Boolean *true, false* to numeric 1,0 respectively.
- I, T are real-valued expressions over the variables of the program, in the interval $[0, 1]$.
- wp and wlp are the probabilistic generalisations of Dijkstra's weakest- and weakest-liberal precondition functions respectively [6, 18, 25, 20].
- G is the loop guard, for us a predicate characterising S_* .
- I is the loop invariant, for us (confusingly) the variant V .
- \Rightarrow is the \leq relation on functions, defined pointwise (as usual).
- $[G]*I \Rightarrow \text{wlp.body}.I$ says that the expected value of I (our V) after one transition is at least as great as its actual value at the state from which the transition was taken. (The $[G]*$ means that the relation is mandated only when G holds, i.e. within S_* .) Thus it is this condition that states that V is a sub-martingale on S_* .
- In general $A \& B$ is $A + B - 1 \max 0$ when $0 \leq A, B \leq 1$. Thus when $T=1$ we have that $I \& T = I$.
- \overline{G} is G 's negation, so that $[\overline{G}]*I$ is our V set to zero on S_* , equivalently our V restricted to S_0 .
- The lemma's inequality thus states
 If termination occurs with probability one ($T=1$ everywhere), then
 the value of V on the starting state is no more than $(I \& T \Rightarrow)$
 the expected value of V on S_0 when escape from S_* has occurred
 ($\text{wp.loop}([\overline{G}]*I)$).